

Northeastern University Runs at the TREC12 Crowdsourcing Track

Maryam Bashir, Jesse Anderton, Jie Wu, Matthew Ekstrand-Abueg,

Peter B. Golbus, Virgil Pavlu, Javed A. Aslam

College of Computer and Information Science

Northeastern University, Boston, USA

{maryam, jesse, evawujie, mattea, pgolbus, vip, jaa}@ccs.neu.edu

February 6, 2013

Abstract

The goal of the TREC 2012 Crowdsourcing Track was to evaluate approaches to crowdsourcing high quality relevance judgments for images and text documents. This paper describes our submission to the Text Relevance Assessing Task. We explored three different approaches for obtaining relevance judgments. Our first two approaches are based on collecting a limited number of preference judgments from Amazon Mechanical Turk workers. These preferences are then extended to relevance judgments through the use of expectation maximization and the Elo rating system. Our third approach is based on our Nugget-based evaluation paradigm.

1 Introduction

We have participated in the Text Relevance Assessing Task (TRAT), one of the two TREC 2012 Crowdsourcing Track tasks. The goal of the TREC Crowdsourcing Track was to evaluate approaches to crowdsourcing high quality relevance judgments for text documents and images. We used the following approaches for obtaining relevance judgments:

- **pair-based ELO** The Elo rating system is a method for calculating the relative rating of players in two player games like Chess [5]. We have used this system for obtaining relevance ranking of documents instead of players by treating a preference judgment between two documents as a match between two players. The Elo rating system updates the score of a document when a certain number of preference comparisons for that document have been made. The score of every document is updated based on the score of the documents to which it was compared. In this work, we demonstrate that the Elo rating system can accurately rank documents using only $O(n)$ preference judgments.
- **pair-based EM** In this context, the Expectation Maximization algorithm [7] is a means of estimating “true” labels from crowd workers as latent variables in a model of worker quality. We convert the collected preferences into binary relevance judgments, and use EM to compute probabilities of relevance for each document.
- **nugget-based** Relevant nuggets—atomic pieces of factual information—are extracted by assessors directly from documents. The nuggets are used to find relevant documents, using a technique based on our nugget-based evaluation framework [9]. This does not use the document pair preference assessments obtained from Mechanical Turk.

Report Documentation Page			Form Approved OMB No. 0704-0188		
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>					
1. REPORT DATE NOV 2012	2. REPORT TYPE	3. DATES COVERED 00-00-2012 to 00-00-2012			
4. TITLE AND SUBTITLE Northeastern University Runs at the TREC12 Crowdsourcing Track				5a. CONTRACT NUMBER	5b. GRANT NUMBER
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	5e. TASK NUMBER
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Northeastern University, College of Computer and Information Science, Boston, MA, 02115				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Presented at the Twenty-First Text REtrieval Conference (TREC 2012) held in Gaithersburg, Maryland, November 6-9, 2012. The conference was co-sponsored by the National Institute of Standards and Technology (NIST) the Defense Advanced Research Projects Agency (DARPA) and the Advanced Research and Development Activity (ARDA). U.S. Government or Federal Rights License					
14. ABSTRACT The goal of the TREC 2012 Crowdsourcing Track was to evaluate approaches to crowdsourcing high quality relevance judgments for images and text documents. This paper describes our submission to the Text Relevance Assessing Task. We explored three different approaches for obtaining relevance judgments. Our first two approaches are based on collecting a limited number of preference judgments from Amazon Mechanical Turk workers. These preferences are then extended to relevance judgments through the use of expectation maximization and the Elo rating system. Our third approach is based on our Nugget-based evaluation paradigm.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1.1 Preference Judgments vs. Relevance Judgments

Traditionally, assessors are asked to give absolute relevance grades to each document with respect to some topic. However, studies have shown that assessors can give more reliable judgments if they are asked which of a pair of documents they prefer, i.e. “is document A better than document B?” rather than “how relevant is document A?” [4]. Another advantage of using preferences is that many popular learning-to-rank algorithms such as RankBoost [6] and RankNet [3] are trained on preferences. Since preferences are often not available, these algorithms need to infer preferences from the absolute relevance judgments collected from assessors. Some information is lost during this process, leading to many ties between documents. This suggests that preferences can be used to improve the training of learning-to-rank algorithms that use such a pairwise approach.

The use of preference judgments on document pairs, as opposed to absolute judgments on documents, for IR evaluations creates new challenges. There are $\binom{n}{2}$ unique pairs of documents for a list of n documents, which means that the number of judgments we need to collect increases to $O(n^2)$. Since collecting judgments is very costly, we need a mechanism for collecting these preferences efficiently. If the collected preferences were transitive, we could reduce the cost by collecting preference judgments for only $O(n \lg n)$ pairs of documents and using them as the comparator for a sorting algorithm. This would produce a sorted list of documents. Unfortunately, preference judgments are known to be not perfectly transitive. We overcome this problem by using the Elo ranking system to produce ranked lists of documents from incomplete preferences.

1.2 Relevant Nuggets vs Relevance Judgments

In a separate run, we propose a new method for relevance assessment based on relevant *information*, not relevant *documents*. Once the relevant information is collected, any document can be assessed/inferred for relevance [9]. The problem with document judgments and with its variants is that the information relevant to a topic is encoded by *documents*, and in the presence of large topic sets or large and/or dynamic collections, it is difficult or impossible to find and judge all relevant documents. Our thesis is that while the number of *documents* potentially relevant to a topic can be enormous, the amount of *information* relevant to a topic, the nuggets, is far, far smaller.

The fundamental thesis of this idea is that finding relevant information immediately leads to relevant documents, which in turn lead to more relevant information. Additionally, documents marked non-relevant are useful for determining non-relevant information; in particular, the TREC assessment philosophy that a document is relevant if it contains any relevant information fits this setup well: any nuggets found in a document marked non-relevant *must be non-relevant*, up to reasonable assessor disagreement and label noise.

The rest of paper is organized as follows: Section 2 describes the collection of preference judgments on document pairs, providing details on our HIT design, experimental setup, and results. Section 3 describes the *pair-based* runs that use the Elo rating system and EM approaches. Section 4 describes our work using the Nugget-based evaluation paradigm. The final section, Section 6, concludes the paper with a description of future work.

2 Document-Pair assessment collection with Mechanical Turk

In this section we describe our experimental design for the collection of preference judgements from crowd workers and our methodology for obtaining complete relevance judgments from only $O(n)$ preference judgements. This section is organized as follows: Section 2.1 gives the details of our Amazon Mechanical Turk setup for crowdsourcing preference judgements, section 2.1.1 describes our HIT design, section 3.1 and section 5 give the methodology and results of using the Elo rating method for extracting relevance judgements from incomplete preference pairs, and section 3.2 gives the methodology and results of using EM for extracting relevance judgements from incomplete preference pairs.

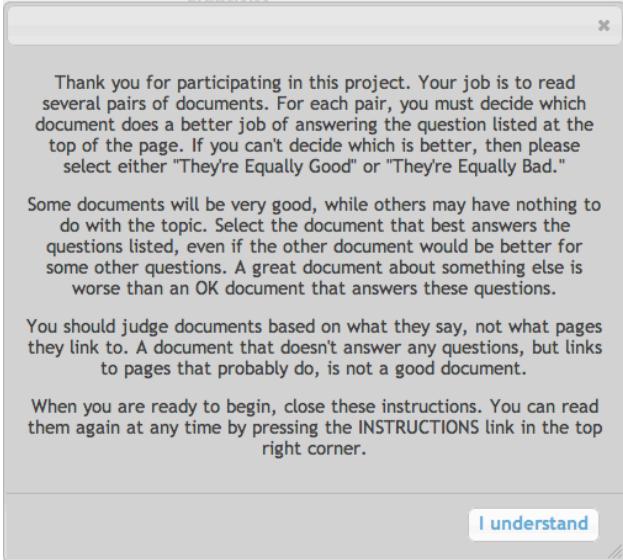


Figure 1: Instructions for preference judgements for crowd workers

2.1 Amazon Mechanical Turk HITs

Amazon Mechanical Turk (AMT) is a crowdsourcing Internet marketplace that gives developers the ability to use human intelligence for tasks that computers are currently unable to do [1]. A person or organization, termed the “requestor,” creates a task definition in the marketplace which workers may then carry out in return for payment. A task definition is called a Human Intelligence Task (HIT). When a worker submits work for a given HIT, the quality of the submission can be automatically assessed and the submission can be either approved, leading to payment, or rejected. There are around 200,000 registered workers of AMT from almost 100 countries. More details on how to use this service can be found in the developer documentation [1].

2.1.1 HIT Design

The HITs we created to collect preference judgements from crowd workers had the following design. After accepting the assignment, workers were shown the instructions seen in Figure 1. In these instructions, we explained that documents should be preferred strictly based on whether they provide information about the query, description, and narrative for a particular topic: that a well-written discussion of a related topic should not be preferred to a poorly-written document which is exactly on topic.

After dismissing the instructions, the worker is shown the interface presented in Figure 2. The “Title” field of a TREC topic is displayed on top, along with its description and narrative. This information describes in detail what constitutes a relevant document for this query. Below the query information is a series of buttons, which allow workers to record their preferences. Two documents are displayed side-by-side below the buttons. The leftmost and rightmost buttons are labeled “This One,” with an arrow pointing to the left or right document, respectively. These buttons allow users to choose a winning document. Between these buttons are two buttons for recording ties, labeled “They’re Equally Good” and “They’re Equally Bad.” Each HIT consisted of 20 preference pairs for the same topic, and had a time limit of 30 minutes. Workers were paid \$0.15 for each approved HIT. The order in which the preferences were displayed, as well as the order of the documents for each particular preference, was randomized.

Document pairs were selected for judgement in the following manner. First, we calculated a prior relevance score for each document using BM25. This produced an initial ranking of the documents for each topic. For each document below the top six, we selected five documents to be compared against, uniformly at random,

Please select the document which provides better information on the topic:

creativity

Questions: Find ways of measuring creativity.
Relevant items include definitions of creativity, descriptions of characteristics associated with creativity, and factors linked to creativity.

<input type="button" value="+ This One"/>	<input checked="" type="button" value="★ They're Equally Good ★"/>	<input type="button" value="✗ They're Equally Bad ✗"/>	<input type="button" value="This One +"/>
<p>AGENCY: National Institute of Standards and Technology Commerce.</p> <p>SUMMARY: The inventions listed below are owned by the U.S. Government, as represented by the Department of Commerce, and are available for licensing in accordance with 35 U.S.C. 207 and 37 CFR Part 404 to achieve expeditious commercialization of results of federally funded research and development. FOR FURTHER INFORMATION CONTACT: Technical and licensing information on these inventions may be obtained by writing to: Marcia Salkeld, National Institute of Standards and Technology, Office of Technology Commercialization, Physics Building, Room B-&-256, Gaithersburg, MD 20899; Fax 301-435-8694; 2751. Any request for information should include the NIST Docket No. and Title for the relevant invention as indicated below.</p> <p>SUPPLEMENTARY INFORMATION: The inventions available for licensing area: NIST Docket No. 90-030D Title: Monomers For Double Ring-Opening Polymerization With Expansion Description: NIST researchers have created a new class of monomers that undergo double ring-opening polymerization with an expansion in volume. When used in resinous compositions the result is a volume neutral curing process at ambient temperature, and a final product that exhibits high adhesive strength. NIST Docket No. 90-036 Title: Epitaxial Iron Films Exhibiting Large Polar Kerr Rotation Description: The invention is a magneto-optic iron film that greatly enhances Kerr rotation compared with conventional iron films. The material could be utilized in magneto-optic data storage media. NIST Docket No. 93-028C</p>			
<p>her father's death when she was 8; by attempted suicide and four months of treatment in a mental hospital while she was in college; by her husband's faithlessness, both imagined and real, which caused her to end their marriage; and finally by her suicide on Feb. 11, 1963.</p> <p>In trying to account for an ending one can only perceive as tragic, Stevenson has chosen to read Plath's life in conventional psychological terms, although her "diagnosis" remains murky. Repeated references to "divided being" recall R. D. Laing's work on schizophrenia, but just as often she mentions "depressed" and "manic" extremes, suggesting bipolar depression, a diagnosis apparently supported by the doctor who was treating Plath when she ended her life. Stevenson makes this bipolarity her controlling metaphor. As early as high school, "Sylvia had a rare, infectious capacity for exultation -- as great a gift for rapture as she had for misery." Indeed, in Stevenson's view, there were two Sylvia Plaths: "the outer Sylvia, characterized by Robert Lowell as 'a brilliant tense presence, embarrassed by restraint,' and the inner woman, fraught with fears and aggressions." And although at one point Stevenson claims that "the writer was beginning to identify with the woman, the woman with the writer; there could be no true distinction between them," her later evaluation is more sinister: that Plath had projected "the 'desired image' (the required image) of herself as Eve -- wife, mother, homemaker, protector of the wholesome, the good, and the holy, an identity that both her upbringing and her own instinctive physical being had fiercely aspired to. Now her submerged and subversive self, utterly true to itself, utterly detached, completely the artist, turned on the Eve scenario and</p>			

Figure 2: Preference pair selection interface with topic keywords and description

from the set of documents with higher BM25 scores. For the top six documents, we collected complete pairwise preferences. We collected four worker assessments for each preference pair which we selected for judgement.

2.1.2 Quality Control

The workers we employed have no particular training in assessing document relevance, so we need a means of verifying the quality of their work. There have been many studies on assessing worker quality for crowdsourcing platforms such as AMT [7]. We used trap questions in our study to ensure that workers are making a reasonable effort, instead of clicking randomly.

A trap question is a question inserted into the HIT which looks the same as any other, but for which a correct answer is known ahead of time. We asked five IR graduate students to create our trap questions by pairing documents which appeared to be highly relevant with documents which appeared to be highly non-relevant. We then mixed five of these trap questions, selected at random, into each HIT. As a result, each assignment consisted of five trap questions and fifteen “real” questions. A worker’s submission was accepted if at least two of the five trap questions were answered correctly, and rejected otherwise.

2.1.3 Data

There were a total of 18,260 documents from TREC 8 ad-hoc, which used the Text Research Collection Volumes 4 (May 1996) and 5 (April 1997) minus the Congressional Record (CR). 10 topics were selected randomly from TREC topics. Participants of the crowdsourcing track were required to simulate the role of NIST assessors for the 10 TREC 8 ad-hoc topics.

3 Pair-based runs

3.1 The Elo Rating System

The Elo rating system is a method for calculating the relative rating of players in two player games [5]. The rating system assigns each player a rating score, with a higher number indicating a better player. Each player's rating is updated after he or she has played a certain number of matches, increasing or decreasing in value depending on whether the player won or lost each match, and on the ratings of both players competing in each match—beating a highly rated player increases one's rating more than beating a player with a low rating, while losing to a player with a low rating decreases one's score more than losing to a player with a high rating. These scores are used in two ways: 1) players are ranked by their scores, 2) the scores can be used to compute the likelihood of each player beating every other player. If the matches are selected intelligently, this can be accomplished even if only $O(n)$ matches are played.

For our problem, we treat each document as a player and each preference judgment between two documents as a match between players. All documents enter the “tournament” rated equally. After each document “plays” a certain number of matches, we update each document’s rating according to equation 3. After all the matches are played, we can rank the documents by their final score. This list can be thresholded to produce absolute relevance judgments. We can also use the scores to compute transitive preference judgments.

For our preliminary experiments, we select $O(n)$ matches stochastically. We wish to sample pairs in such a way that we create a bias towards relevant documents. In this way, relevant documents will play more matches than non-relevant documents, giving them more opportunities to improve their ratings and move up the list. We begin by ranking documents for each topic using the BM25 retrieval method. The first five documents all “play” one another. Each remaining document in the list plays against 5 randomly selected documents with higher ranks. In our experiments, every document plays at least 5 matches. On average, every document plays 11 matches. We sort all documents based on their ratings after all $O(n)$ matches have been played.

3.1.1 Mathematical Details of The Elo Rating System

The Elo rating system measures the skill level of players in relative terms. Each player’s rating depends on the rating of his or her opponents, as well as the of wins or losses. The expected score for each player in a match can be estimated from the ratings of the players. If player A has rating R_A and player B has rating R_B , then the expected score of player A is calculated as follows:

$$E_A = \frac{1}{1 + 10^{\frac{R_B - R_A}{200}}} \quad (1)$$

Similarly the expected score for player B is:

$$E_B = \frac{1}{1 + 10^{\frac{R_A - R_B}{200}}} \quad (2)$$

Player’s ratings are updated by comparing their actual score to their expected score. When a player performs worse than expected, then his or her rating is decreased. If player A has an expected score, E_A , and an actual score, S_A , then his or her rating will be updated by following formula:

$$R'_A = R_A + K(S_A - E_A) \quad (3)$$

This update can be made after every match, or after several matches according to the situation.

ELO Experimental Setup. We have set K equal to 20 in Elo rating update equation 3 and each document has an initial uniform rating of 1000. Every document pair selected for comparison was judged by 5 crowd

workers. For each of the 10 topics, participants of TREC Crowdsourcing track were required to provide both binary relevance judgments and probabilities of relevance for each document. Probabilities were produced by normalizing the Elo scores. We then label the n most probable documents as relevant, where n was chosen by manual inspection.

3.1.2 TREC Runs “NEUElo2/3/4/5”

We have submitted four different runs to the TREC Crowdsourcing track using the Elo rating algorithm for preference judgments. Due to time constraints we were not able to collect five judgments for every document pair for our submitted runs to TREC. The following is a description of each of the runs created using the Elo rating algorithm:

NEUElo2: This run was created using a non-uniform number of judgments for each document pair, i.e. the number of judgments for each document pair could be from 1 to 5. The threshold for binary relevance was the top 20 documents for each topic.

NEUElo3: This run was the same as NEUElo3 except that we also used the judgment from our trap questions. Since the trap questions were created by experts, we added the documents preferred in the trap questions to the top of each ranked list. We also chose a different number of relevant documents for topic, again by manual inspection.

NEUElo4: This run was the same as NEUElo3 except that we ignored all the document pairs that had only 1 judgment. This was done to avoid random judgments.

NEUElo5: This run was the same as NEUElo4 except that the threshold for binary relevance was the top 20 documents for each topic.

3.2 Expectation Maximization run “NEUEM1”

Expectation Maximization (EM) is an iterative method for finding the maximum likelihood estimate of the parameters of a probability distribution using a model incorporating unobserved latent variables. In our model, we treat the observed preference judgments from Crowdsource workers as being drawn from a distribution parameterized by the “quality” of each worker. The relevance grade of each document is treated as a latent variable. For the purposes of this task, the end result is a probability distribution over each latent variable, i.e. the probability that each document is relevant. Our approach is similar to that of Hosseini et al. [8].

3.2.1 The EM Algorithm

Assume that there are N documents and M crowd workers. Let G represent the number of relevance grades. We assign each worker j a $G \times G$ latent confusion matrix, C^j . Each element c_{lg}^j is the probability that worker j assigns a label l to a document whose true relevance is g . In our model, each document has one of two relevance grades, relevant and non-relevant; i.e. $g \in \{r, n\}$. Hence C^j is a 2×2 confusion matrix, $C^j = \begin{pmatrix} c_{rr} & c_{rn} \\ c_{nr} & c_{nn} \end{pmatrix}$. Let $\Pr(R_i = g)$ denote the probability that document i has the true relevance grade g , p_g

is the prior probability distribution over relevance grades, and X_{il}^j is an indicator variable indicating whether worker j has assigned labeled l to document i .

There are five steps in the EM algorithm:

Step 0: Pre-processing

We wish to learn the relevance grades of each document. However, we have collected preferences between pairs of documents. Before we can use the EM algorithm, we must first convert these preferences into grades. We do so in the following way: assume that worker j prefers document A to document B . Then we interpret this as worker j labeling document A as relevant, and document B as non-relevant.

Step 1: Initialization

We assume that all of the crowd workers are honestly attempting to give correct answers. Therefore, we initialize each workers confusion matrix to:

$$C_j = \begin{vmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{vmatrix}$$

Furthermore, we assume that all documents are equally likely to be relevant and non-relevant. Therefore, we initialize $p_r = 0.5$ and $p_n = 0.5$.

Step 2: Estimate the probability of relevance

In this step, we combine the confusion matrix and relevance labels from each worker to estimate the probability that each document is relevant. All labels are assumed to be i.i.d.

$$\Pr(R_i = r \mid C_j, \forall j \in M) = \frac{p_r \times \prod_{j=1}^M \prod_{l=0}^G c_{lr}^j X_{il}^j}{p_r \times \left(\prod_{j=1}^M \prod_{l=0}^G c_{lr}^j X_{il}^j + \prod_{j=1}^M \prod_{l=0}^G c_{ln}^j X_{il}^j \right)} \quad (4)$$

$$\Pr(R_i = n \mid C_j, \forall j \in M) = \frac{p_r \times \prod_{j=1}^M \prod_{l=0}^G c_{ln}^j X_{il}^j}{p_r \times \left(\prod_{j=1}^M \prod_{l=0}^G c_{lr}^j X_{il}^j + \prod_{j=1}^M \prod_{l=0}^G c_{ln}^j X_{il}^j \right)} \quad (5)$$

Step 3: Estimate the maximum likelihood

Using the each documents' estimated probability of relevance calculated in step 2, we produce an updated estimate of each worker's confusion matrix.

$$c_{lg}^j = \frac{\sum_{i=1}^N \left(\Pr(R_i = g) \times X_{il}^j \right)}{\sum_{k=0}^G \sum_{i=i}^N \left(\Pr(R_i = g) \times X_{ik}^j \right)} \quad (6)$$

Step 4: Iterate Step 2 and 3 until convergance

We repeatedly calculated worker confusion matrices and probabilities of document relevance until the results converged. We consider this process to have converged when, for each document i , the difference between $\Pr(R_i = g)$ at iteration $t - 1$ and at iteration t for all documents i and relevance grades g is less than or equal to 0.01.

After applying the steps above, we created ranked lists by sorted the documents by their probability of relevance, i.e. $\Pr(R_i = r)$. Additionally, all documents whose probability of relevance is greater than 0.5 is marked as relevant.

4 Nuggets-based run “NEUNugget12”

Separate from pair assessments runs using mechanical turk, we produced a run using our nugget technology [9]. We use a framework of mutual, iterative reinforcement between nuggets and documents, most similar to Active Learning [10, 2] mechanisms. The human feedback (relevance assessment) is given *iteratively* on the documents/data; thus, while the assessor judges the document (as it is standard IR procedure), our “back-end” system infers the relevance of unjudged documents as well as the relevance of text fragments, or nuggets, from within the documents.

Figure 3 illustrates this framework, with MATCHING being the reinforcing procedure from documents to nuggets and vice-versa:

- the iterative loop selects documents based on the current notion of relevance (as expressed by current set of nuggets GREL); documents are judged and put into the set of documents QREL;
- for relevant documents, nuggets are extracted and added to GREL
- all nugget scores are updated based on the relevant and non-relevant documents judged so far and how well each nugget matches them

This framework uses four components that can be designed somewhat independently. Our implementa-

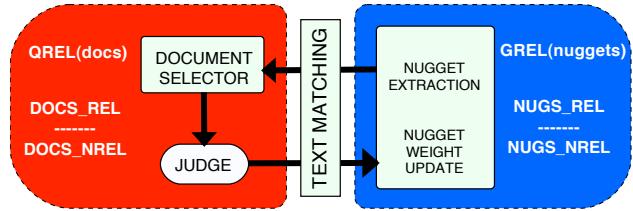


Figure 3: The overall design of assessment process: Iteratively, documents are selected and assessed, and nuggets are extracted and [re]weighted.

Nugget Extractor

Document to be Judged

✓ Relevant ✘ Nonrelevant ⌂ Refresh Reload

Thursday February 3, 1994 Part III Environmental Protection Agency 40 CFR Part 763 Asbestos Model Accreditation Plan; Interim Final Rule Federal Register /Vol. 59, No. 23./ Thursday, February 3, 1994./Rules and Regulations ENVIRONMENTAL PROTECTION AGENCY 40 CFR Part 763 [GPO:2001-02-01:2010-02-01:FR-2001-02-01:2001-02-01] RIN 2070-Abyp;AC51 AGENCY: Environmental Protection Agency (EPA). ACTION: Interim final rule.

ENVIRONMENTAL PROTECTION AGENCY (EPA) is issuing an interim final rule to revise its asbestos Model Accreditation Plan (MAP) to clarify the types of persons who must be accredited to work with asbestos in schools and public and commercial buildings; to increase the minimum number of hours of training, including additional hours of hands-on health and safety training for asbestos abatement contractors; to reflect changes to asbestos abatement standards and other necessary changes as mandated by section 15(a)(3) of the Asbestos School Hazard Abatement Reauthorization Act (ASHARA). This revised rule replaces the original MAP found at 40 CFR part 763, Appendix A. Appendix A contains the revised MAP, which, when taken together, comprised a model asbestos accreditation plan for States and EPA-approved training providers. These components included: (1) initial training, (2) reresher training, (4) qualification, (5) decertification requirements, and (6) recordkeeping requirements. The revised MAP includes new map definitions, which help to determine the scope and applicability of the rule, and (2) new recordkeeping requirements for the providers of accredited training courses. The changes also specify the deadline for States to modify their accreditation programs to be no less stringent than the revised MAP. EPA is also amending the revised MAP to add section 206(b) (2). Further, the revised MAP prescribes deadlines for training course providers and persons who must obtain accreditation to comply with new requirements; distinguishes between the training requirements for the revised accreditation training courses; adds several new topics to the project designer training curriculum; establishes new enforcement criteria and Federal procedures for withdrawing approval from accredited persons and training programs; and stipulates new information requirements for training course providers. In addition, the revised MAP specifies that each state's accreditation plan, States may have to modify their programs to insure that each state has a contractor accreditation plan that is at least as stringent as the revised MAP as required by TSCA section 206(e). Similarly, training providers may need to adjust their training courses and training programs to comply with the revised MAP. EPA will modify the organization, and some of the language of the original MAP. These modifications, however, are technical, and do not impose new substantive requirements.

DATES: This rule is effective April 4, 1994. While this is an interim final rule, EPA is accepting further comments on this action. All written comments must be received by EPA no later than March 4, 1994. EPA will consider the written comments received during the 30-day comment period in determining the need for any further rule amendments.

EXPLANATION: Written comments are invited to the Office of Pollution Prevention and Toxics Management Division (7404), Office of Pollution Prevention and Toxics (OPPT), Environmental Protection Agency, 401 M St., Washington, DC 20460. EPA does not anticipate receiving any comments that contain information claimed as confidential business information (CBI). If such information is contained in any comments, it must be clearly marked as CBI. Any information claimed as CBI or they will be placed in the public record. CBI claims should be accompanied by statements substantiating the claim as described in 40 CFR 2.204(e)(4). If information is claimed as CBI, a nonconfidential version of the comments should also be submitted for review by the agency.

FOR FURTHER INFORMATION CONTACT: Susan B. Hazen, Director, Environmental Assistance Division (7408), Office of Pollution Prevention and Toxics, Environmental Protection Agency, 401 M St., Washington, DC 20460. EPA does not anticipate receiving any comments that contain information claimed as confidential business information (CBI). If such information is contained in any comments, it must be clearly marked as CBI. Any information claimed as CBI or they will be placed in the public record. CBI claims should be accompanied by statements substantiating the claim as described in 40 CFR 2.204(e)(4). If information is claimed as CBI, a nonconfidential version of the comments should also be submitted for review by the agency.

© 2012 Northeastern University

DB: treccs_nuggets Query: Three Gorges Project Go!

Unjudged Nuggets	
Mark	Nugget + Add Nugget ⌂ Refresh
✓	These approaches are being studied as a part of the planning work. Normal water level should be selected primarily based o... 1...
✓	The other 14 units shall mainly be manufactured at home, with the Chinese plants becoming major suppliers of the goods. I... 1
✓	Areas along the recently completed line, which have long been a hub of telecommunications, previously relied on intermedi... 1
✓	The 4,700-km line, which passes through: eight provinces and municipalities and 22 prefectures and cities, links Beijing wit... 1
✓	Since that time the project has maintained an excellent safety record and has operated smoothly. 1
✓	Preparations are also being made to begin construction of Huanghua Port, a special coal dock with an annual handling cap... 1
✓	The coalfield, located in north Shaanxi Province and the southern part of the Inner Mongolian Autonomous Region, has veri... 1
✓	The project is designed to cater to the strategic plan calling for a westward shift in the nation's energy foundation. 1
✓	The complexity of the technology and scale of construction are unprecedented in China. The 21 billion yuan project, funde... 1
✓	Ertan, which features China's first 240-meter-high dam, will also have the country's largest turbo-generators with a generat... 1
✓	The 3,500-km expressway, which will link Qinhuangdao, Lianyungang and other important port cities in east China's coastal... 1
✓	The 2,300-km north-south thoroughfare will pass through major cities in Hebei, Henan, Hubei, Hunan and Guangdong prov... 1
✓	The 850-km expressway, a section of the Tokyo-Pyongyang-Seoul-Beijing-Moscow-London international highway transport ... 1

Relevant Nuggets	
Mark	Nugget + Add Nugget ⌂ Refresh
✗	By the end of this century, China should have the capability to manufacture 800MW generator units. It is also possible to jo... 1...
✗	The massive project, which calls for a 20-year construction period, is expected to cost 33 billion yuan, based on constant p... 1
✗	Article Type:BN [Text] Yichang, February 27 (XINHUA) -- The construction of China's largest water conservancy scheme, th... 1
✗	It will affect about 12 million residents in the middle reaches of the Yangtze and take 17 years to finish. Construction starte... 1
✗	As a result, it will slow down the rise of flood levels at the end of the Three Gorges caused by silting. (3) According to an a... 1
✗	A total of 24 units should be installed with a total capacity of 19.20 MW; 1.52 MW higher than that in the validation plan.... 1
✗	The total installed capacity is 700MW x 26 = 18,200MW, which is 520,000 kW higher than that in the validated plan. Seco... 1
✗	The output of the turbine is 729 MW and the generator unit produces 700 MW. 1
✗	Installed Capacity as a Function of Reduced Operating Hours at t The Three Gorges Hydropower Station 1
✗	Yearly Operating Hours of Power Plants 1 MW or Larger Already Implied and Under Construction in China 1
✗	It should match the total flow passing through all generators up and down the river. Based on the superior central geograph... 1
✗	Hopefully, this will not happen at the Three Gorges. The plan to install 24 800MW units at the Three Gorges is of vital impo... 0...
✗	It is recommended that 24 generator units with an individual unit capacity of 800MW be placed. 0...

Nonrelevant Nuggets	
Mark	Nugget + Add Nugget ⌂ Refresh

Figure 4: The nugget extraction interface, on query “Three Gorges Project”, showing a document to be judged (left) and the list of candidate nuggets (right)

tions for SELECTION, MATCHING, and EXTRACTION are presented in [9], noting that for each of these we settled on these methods after trying various ideas. While ours work well, each can be independently replaced by more suited techniques for specific tasks.

A fifth component, not explored in this paper, is the JUDGE model. Here we assume the judge follows the traditional NIST assessor, simply labeling binary documents as relevant (contains something of some relevance) or non-relevant; a more sophisticated judge can review and modify the nuggets extracted, categorize documents, use graded relevance, annotate important keywords, or classify the query.

The supervised aspect of the nugget extraction process is shown in Figure 4. Here the user views the next document to be judged as well as the current relevance beliefs of all extracted nuggets. One aspect of the process added to our previous work is the ability to mark nuggets as relevant and nonrelevant. This is achieved through the buttons next to each nugget on the right side of the interface. While not necessary, if a user judges a nugget, we can integrate this information into our update system.

Every nugget has a score that is a combination of the relevance judgments of all matching judged documents. A matching relevant document increases the score, and a nonrelevant document decreases the score. Therefore, if a nugget is marked relevant, we set its score to the maximum of all unjudged nuggets, and if a nugget is marked nonrelevant, it is given the minimum score of all unjudged nuggets. As described in our previous paper, these scores are then used to update the relevance belief of unjudged documents, and to then select the next documents for judgment. Finally, when the process stops (usually up to the assessor), to obtain binary relevance judgments from document scores, we use a threshold of 0.8.

This resulted in a low amount of manual work per query, as a small percentage of documents were judged and the relevance of the rest were extrapolated. The entire nugget extraction process took about one hour per query and was performed by a grad student. While only one person was used, the methodology could utilize judgments from multiple workers to reduce user error. We performed our experiments with one person to demonstrate the small amount of manual work required to infer relevance grades for the entire set.

5 Results

Following the conventions of the TREC Crowdsourcing track, we evaluate binary preference judgements using both Logistic Average Misclassification (LAM) and the area under the ROC curve (AUC). The evaluation results of our runs submitted to TREC are given in Table 1.

Topic ID	NEUElo2		NEUElo3		NEUElo4		NEUElo5		NEUEM1		NEUNuggets12	
	AUC	LAM	AUC	LAM	AUC	LAM	AUC	LAM	AUC	LAM	AUC	LAM
411	0.665	0.293	0.69	0.11	0.693	0.145	0.692	0.171	0.650	0.383	0.860	0.121
416	0.735	0.265	0.745	0.209	0.854	0.183	0.855	0.162	0.623	0.432	0.868	0.220
417	0.678	0.261	0.695	0.195	0.716	0.182	0.716	0.156	0.728	0.196	0.686	0.132
420	0.579	0.285	0.598	0.134	0.642	0.153	0.639	0.151	0.524	0.464	0.865	0.157
427	0.747	0.146	0.756	0.134	0.788	0.122	0.788	0.12	0.605	0.260	0.798	0.191
432	0.598	0.38	0.598	0.256	0.631	0.38	0.631	0.38	0.624	0.405	0.580	0.375
438	0.578	0.272	0.583	0.38	0.62	0.319	0.617	0.311	0.578	0.154	0.762	0.447
445	0.775	0.232	0.783	0.12	0.706	0.144	0.706	0.212	0.691	0.452	0.891	0.121
446	0.683	0.392	0.72	0.227	0.729	0.238	0.711	0.156	0.656	0.460	0.707	0.274
447	0.78	0.265	0.953	0.012	0.992	0.051	0.953	0.082	0.652	0.186	0.464	0.081
All	0.682	0.279	0.712	0.178	0.737	0.192	0.731	0.19	0.633	0.339	0.748	0.213

Table 1: Evaluation results of runs submitted to TREC 2012 Crowd Sourcing track.

6 Conclusion and Future Work

In this paper we have described our work based on three approaches for obtaining high quality relevance judgements. The first two approaches are based on crowd sourcing relevance judgements using preferences,

whereas the third approach is based on nuggets assessments.

Elo runs use partial preference judgements obtained from crowd workers as input to a popular rating system that outputs a rank list of documents with their relevance scores. Our preliminary results submitted to TREC crowd sourcing track are close to average results submitted to TREC by various systems for most of the topics. We later have experimented with some modifications to our basic Elo rating system, and collected more crowdsourced pair assessments, which resulted in significant improvement to our reported TREC 2012 results. For our future work, we plan to design more intelligent matches between documents, with a goal to obtain high quality judgements from small number of preference judgements.

The EM run is a straight forward application of Expectation Maximization algorithm on crowdsourced assessments, as other researchers have done, in order to account for worker quality. We use it as a baseline comparison, but we think it can be most useful as an intermediary processing stage of crowd workers assessments, before running other methods like Elo. One of the Elo enhancements mentioned above is to run it over EM-processed pair assessments.

The nuggets run is very different than the others, to the point that “crowdsourcing” might not be an appropriate name: the focus here is not on cheap weak judgments that can be averaged somehow into a good assessment, but rather on very few judgments of high quality (on nuggets) and some expert feedback on documents. We note that the nuggets run did best among the runs we submitted; while not a lot better, the human effort put in to it, even at expert level, is significantly less than either crowdsourcing or TREC assessing efforts. Unfortunately, later analysis showed that our “expert” grad student performing the assessments disagreed significantly, perhaps mistakenly, with the existing TREC qrel assessments, thus certainly impacting the run performance.

References

- [1] Amazon mechanical turk. <http://www.mturk.com>.
- [2] Y. Baram, R. El-Yaniv, and K. Luz. Online choice of active learning algorithms. *J. Mach. Learn. Res.*, Volume 5, December 2004.
- [3] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, ICML ’05, pages 89–96, New York, NY, USA, 2005. ACM.
- [4] B. Carterette, P. N. Bennett, D. M. Chickering, and S. T. Dumais. Here or there. In *ECIR*, pages 16–27, 2008.
- [5] A. Elo and S. Sloan. *The Rating of Chess Players, Past and Present*. Arco Publishing, New York, NY, USA, 1978.
- [6] Y. Freund, R. Iyer, R. E. Schapire, and Y. Singer. An efficient boosting algorithm for combining preferences. *J. Mach. Learn. Res.*, 4:933–969, Dec. 2003.
- [7] M. Hosseini, I. J. Cox, N. Milić-Frayling, G. Kazai, and V. Vinay. On aggregating labels from multiple crowd workers to infer relevance of documents. In *Proceedings of the 34th European conference on Advances in Information Retrieval*, ECIR’12, pages 182–194, Berlin, Heidelberg, 2012. Springer-Verlag.
- [8] G. Kazai, J. Kamps, M. Koolen, and N. Milic-Frayling. Crowdsourcing for book search evaluation: impact of hit design on comparative system ranking. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, SIGIR ’11, pages 205–214, New York, NY, USA, 2011. ACM.
- [9] S. Rajput, M. Ekstrand-Abueg, V. Pavlu, and J. A. Aslam. Constructing test collections by inferring document relevance via extracted relevant information. In *Proceedings of the 21st ACM Conference on Information and Knowledge Management*, pages 145–154. ACM Press, October 2012.

- [10] H. S. Seung, M. Opper, and H. Sompolinsky. Query by committee. In *Proceedings of the Fifth Annual ACM Conference on Computational Learning Theory*, COLT '92, 2002.